Information and Computation ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

A semantic theory of the Internet of Things ☆

Ruggero Lanotte [a], Massimo Merro [b,*]

[a] Dipartimento di Scienza e Alta Tecnologia, Università degli Studi dell'Insubria, Via Valleggio 11, 22100 Como, Italy
[b] Dipartimento di Informatica, Università degli Studi di Verona, Strada le Grazie 15, 37134 Verona, Italy

ARTICLE INFO

Article history:
Received 29 March 2017
Received in revised form 2 December 2017
Available online xxxx

Keywords:
Internet of Things
Process calculus
Operational semantics
Behavioural semantics
Bisimulation

ABSTRACT

We propose a process calculus for modelling and reasoning on systems in the *Internet of Things* paradigm. Our systems interact both with the physical environment, via *sensors* and *actuators*, and with *smart devices*, via short-range and Internet channels. The calculus is equipped with a standard notion of labelled *bisimilarity* which is proved to be a coinductive characterisation of a well-known contextual equivalence. We use our semantic proof-methods to prove run-time properties of a non-trivial case study as well as system equalities.

1. Introduction

In the *Internet of Things* (IoT) paradigm, *smart devices* equipped with embedded technology automatically collect information from shared resources (e.g. Internet accesses, physical devices, etc.) and aggregate them to provide new services to end users [2]. The "things" commonly deployed in IoT systems are: *RFID tags*, for unique identification, *sensors*, to detect physical changes in the environment, and *actuators*, to pass information to the environment. To provide proper communication capabilities, smart devices are organised in networks which are based on the standard communication protocols of the Internet framework.

The range of IoT applications is rapidly increasing and already covers several domains [3,2,4]: (i) environmental monitoring, (ii) healthcare, (iii) personal and social, (iv) security and surveillance, (v) smart environment (home, offices, cities), (vi) transportation and logistics (automotive).

The research on IoT is currently focusing on practical applications such as the development of enabling technologies [5], ad hoc architectures [6], semantic web technologies [7], and cloud computing [2]. However, as pointed out by Lanese et al. [8], there is a lack of research in formal methodologies to model the interactions among system components, and to verify the correctness of the network deployment before its implementation.

The main goal of the current paper is to propose a new process calculus for IoT systems which supports a clear semantic theory for specifying and reasoning on IoT applications. Devising a calculus for modelling a new paradigm requires understanding and distilling, in a clean algebraic setting, the basic features of the paradigm. In order to point out the main ingredients of the IoT paradigm, we use a small example within the smart environment domain.

☆ An extended abstract appeared in the proceedings of the *8th International Conference on Coordination Models and Languages* (COORDINATION 2016), volume 9686 of *Lecture Notes in Computer Science*, pp. 157–174, Springer, 2016 [1].

* Corresponding author.
E-mail address: massimo.merro@univr.it (M. Merro).

https://doi.org/10.1016/j.ic.2018.01.001
0890-5401/

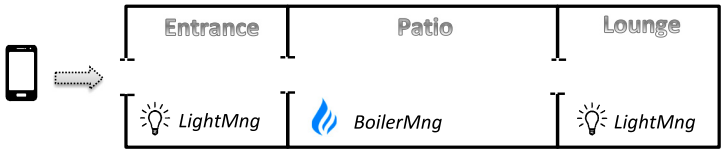

Fig. 1. A simple smart home.

Suppose a simple *smart home* (see Fig. 1) in which the user can (i) profit of her smartphone to remotely control the heating boiler of her house, and (ii) automatically turn on lights when entering a room. The house consists of an entrance and a lounge, separated by a patio. Entrance and lounge have their own lights (actuators) which are governed by different light manager processes, *LightMng*. The boiler is placed in the patio and it is governed by a boiler manager process, *BoilerMng*. This process senses the local temperature (via a sensor) and decides whether the boiler should be turned on/off, setting a proper actuator to signal the state of the boiler.

The smartphone executes two concurrent processes: *BoilerCtrl* and *LightCtrl*. The first one reads user's commands, submitted via the phone touchscreen (a sensor), and forwards them to the process *BoilerMng* of the house, via an Internet channel. Whereas, the process *LightCtrl* interacts with the processes *LightMng* of the house, via short-range wireless channels (e.g. Bluetooth, infrared, etc), to automatically turn on lights when the smartphone physically enters either the entrance or the lounge. The whole system is given by the parallel composition of the smartphone (a mobile device) and the smart home (a stationary entity).

On this kind of systems one may wish to prove interesting *run-time properties*. Think of a *fairness property* saying that the boiler will be eventually turned on/off whenever specific conditions are satisfied. Or *consistency properties*, saying, for instance, that the smartphone will never be in two rooms at the same time. Even more, one may be interested in understanding whether different implementations of our smart home have the same *observable behaviour*. Consider a variant of our smart home, where lights functionality depends on the GPS coordinates of the smartphone (localisation is a common feature of today smartphones). Intuitively, the smartphone could send its GPS position to a centralised light manager, *CLightMng* (possibly placed in the patio), via an Internet channel. The process *CLightMng* will then interact (via short-range channels) with the local light manager processes to turn on/off lights, depending on the current position of the smartphone. Here comes an interesting question: can these two implementations of the smart home, based on different light management mechanisms, be actually distinguished by an end user?

In the paper at hand we develop a fully abstract semantic theory for a process calculus of IoT systems, called `CaIT`. We provide a formal notion of when two systems in `CaIT` are indistinguishable, in all possible contexts, from the point of view of the end user. Formally, we adopt the approach of [9,10], often called *reduction (closed) barbed congruence*, which relies on two crucial concepts: a *reduction semantics* to describe system computations, and *basic observables* to represent what the environment can directly observe of a system. As IoT systems are essentially *cyber-physical systems* [11], they have at least two possible observables: the ability to transmit along channels, *logical observation*, and the capability to modify actuators, *physical observation*. In `CaIT`, we have adopted the second form of observable as our contextual equality remains invariant when adding logical observation. However, the right definition of physical observation is far from obvious as it has a non-trivial impact on the definition of the reduction semantics. Thus, observables and reduction semantics contain *key design choices* for the formal definition of `CaIT`.

Our calculus is equipped with two *labelled transition semantics* (LTSs) in the SOS style of Plotkin [12]: an *intensional* semantics and an *extensional* semantics. The adjective intensional is used to stress the fact that the actions here correspond to activities which can be performed by a system in isolation, without any interaction with the external environment. On the other hand, the extensional semantics focuses on those activities which require a contribution of the environment. Our extensional LTS builds on the intensional one, by introducing specific transitions for modelling all interactions with the environment. Here, we would like to point out that since our basic observation on systems does not involve the recording of the passage of time, this has to be taken into account extensionally.

We prove that the reduction semantics coincides with the intensional semantics (Harmony theorem), and that is satisfies some desirable time properties such as (a localised variant of) *time determinism*, *patience*, *maximal progress* and *well-timedness* [13]. However, the main result of the paper is that *weak bisimilarity* in the extensional LTS provides a *coinductive characterisation* of our contextual equivalence, reduction barbed congruence: two systems are related by some bisimulation in the extensional LTS if and only if they are reduction barbed congruent. Full abstraction results of this kind are in general hard to achieve. In our case, this result required a non-standard proof of the congruence theorem for the weak bisimilarity.

We finally show the effectiveness of our bisimulation proof-technique to deal with non-trivial systems. In particular, we provide a formal proof that two different implementations of the smart home mentioned before are bisimilar. Formal proofs of systems of such size are quite rare in the literature. Thus, in order to reduce the size of the bisimulation relation to be exhibited, we make an intensive use of *up-to expansion* proof-techniques [10].

Outline. Section 2 contains the calculus together with the reduction semantics, the contextual equivalence, and a discussion on design choices. Section 3 gives the details of our smart home example, and proves desirable run-time properties for it. Section 4 defines both intensional and extensional LTSs. In Section 5 we define bisimilarity for (networks of) IoT-systems,

Please cite this article in press as: R. Lanotte, M. Merro, A semantic theory of the Internet of Things, Inf. Comput. (2018), https://doi.org/10.1016/j.ic.2018.01.001